

# PASSCoDe: Parallel ASynchronous Stochastic dual Co-ordinate Descent

Cho-Jui Hsieh  
University of Texas, Austin  
cjhsieh@cs.utexas.edu

Hsiang-Fu Yu  
University of Texas, Austin  
rofuyu@cs.utexas.edu

Inderjit S. Dhillon  
University of Texas, Austin  
inderjit@cs.utexas.edu

## Abstract

Stochastic Dual Coordinate Descent (*DCD*) has become one of the most efficient ways to solve the family of  $\ell_2$ -regularized empirical risk minimization problems, including linear SVM, logistic regression, and many others. The vanilla implementation of *DCD* is quite slow; however, by maintaining primal variables while updating dual variables, the time complexity of *DCD* can be significantly reduced. Such a strategy forms the core algorithm in the widely-used LIBLINEAR package. In this paper, we parallelize the *DCD* algorithms in LIBLINEAR. In recent research, several synchronized parallel *DCD* algorithms have been proposed, however, they fail to achieve good speedup in the shared memory multi-core setting. In this paper, we propose a family of asynchronous stochastic dual coordinate descent algorithms (*PASSCoDe*). Each thread repeatedly selects a random dual variable and conducts coordinate updates using the primal variables that are stored in the shared memory. We analyze the convergence properties when different locking/atomic mechanisms are applied. For implementation with atomic operations, we show linear convergence under mild conditions. For implementation without any atomic operations or locking, we present the first *backward error analysis* for *PASSCoDe* under the multi-core environment, showing that the converged solution is the exact solution for a primal problem with perturbed regularizer. Experimental results show that our methods are much faster than previous parallel coordinate descent solvers.

## 1 Introduction

Given a set of instance-label pairs  $(\dot{x}_i, \dot{y}_i)$ ,  $i = 1, \dots, n$ ,  $\dot{x}_i \in \mathbb{R}^d$ ,  $\dot{y}_i \in \mathbb{R}$ , we focus on the following empirical risk minimization problem with  $\ell_2$ -regularization:

$$\min_{\mathbf{w} \in \mathbb{R}^d} P(\mathbf{w}) := \frac{1}{2} \|\mathbf{w}\|^2 + \sum_{i=1}^n \ell_i(\mathbf{w}^T \mathbf{x}_i), \quad (1)$$

where  $\mathbf{x}_i = \dot{y}_i \dot{x}_i$ ,  $\ell_i(\cdot)$  is the loss function and  $\|\cdot\|$  is the 2-norm. A large class of machine learning problems can be formulated as the above optimization problem. Examples include Support Vector Machines (SVMs), logistic regression, ridge regression, and many others. Problem (1) is usually called the primal problem, and can usually be solved by Stochastic Gradient Descent (SGD) (Zhang, 2004; Shalev-Shwartz et al., 2007), second order methods (Lin et al., 2007), or primal coordinate descent algorithms (Chang et al., 2008; Huang et al., 2009).

Instead of solving the primal problem, another class of algorithms solves the following dual problem of (1):

$$\min_{\boldsymbol{\alpha} \in \mathbb{R}^n} D(\boldsymbol{\alpha}) := \frac{1}{2} \left\| \sum_{i=1}^n \alpha_i \mathbf{x}_i \right\|^2 + \sum_{i=1}^n \ell_i^*(-\alpha_i), \quad (2)$$

where  $\ell_i^*(\cdot)$  is the conjugate of the loss function  $\ell_i(\cdot)$ , defined by  $\ell_i^*(u) = \max_z (zu - \ell_i(z))$ . If we define

$$\mathbf{w}(\boldsymbol{\alpha}) = \sum_{i=1} \alpha_i \mathbf{x}_i, \quad (3)$$

then it is known that  $\mathbf{w}(\boldsymbol{\alpha}^*) = \mathbf{w}^*$  and  $P(\mathbf{w}^*) = -D(\boldsymbol{\alpha}^*)$  where  $\mathbf{w}^*, \boldsymbol{\alpha}^*$  are the optimal primal/dual solutions respectively. Examples include hinge-loss SVM, square hinge SVM and  $\ell_2$ -regularized logistic regression.

Stochastic Dual Coordinate Descent (*DCD*) has become the most widely-used algorithm for solving (2), and it is faster than primal solvers (including SGD) in many large-scale problems. The success of *DCD* is mainly due to the trick of maintaining the primal variables  $\mathbf{w}$  based on the primal-dual relationship (3). By maintaining  $\mathbf{w}$  in memory, Hsieh et al. (2008); Keerthi et al. (2008) showed that the time complexity of each coordinate update can be reduced from  $O(\text{nnz})$  to  $O(\text{nnz}/n)$ , where  $\text{nnz}$  is number of nonzeros in the training dataset. Several *DCD* algorithms for different machine learning problems are currently implemented in LIBLINEAR (Fan et al., 2008) and they are now widely used in both academia and industry. The success of *DCD* has also catalyzed a large body of theoretical studies (Nesterov, 2012; Shalev-Shwartz & Zhang, 2013).

In this paper, we parallelize the *DCD* algorithm in a shared memory multicore system. There are two threads of work on parallel coordinate descent. The first thread focuses on synchronized algorithms, including synchronized CD (Richtárik & Takáč, 2012; Bradley et al., 2011) and synchronized *DCD* algorithms (Yang, 2013; Jaggi et al., 2014). However, choosing the block size is a trade-off problem between communication and convergence speed, so synchronous algorithms usually suffer from slower convergence. To overcome this problem, the other thread of work focuses on asynchronous CD algorithms in multi-core shared memory systems (Liu & Wright, 2014; Liu et al., 2014). However, none of the existing work maintains both the primal and dual variables. As a result, the recent asynchronous CD algorithms end up being much slower than the state-of-the-art serial *DCD* algorithms that maintain both  $\mathbf{w}$  and  $\boldsymbol{\alpha}$ , as in the LIBLINEAR software. This leads to a challenging question: how to maintaining both primal and dual in an asynchronous and efficient way?

In this paper, we propose the first asynchronous dual coordinate descent (*PASSCoDe*) algorithms with the address to the issue for the primal variable maintenance in the shared memory multi-core setting. We carefully discuss and analyze three versions of *PASSCoDe*: *PASSCoDe-Lock*, *PASSCoDe-Atomic*, and *PASSCoDe-Wild*. In *PASSCoDe-Lock*, convergence is always guaranteed but the overhead for locking makes it even slower than serial *DCD*. In *PASSCoDe-Atomic*, the primal-dual relationship (3) is enforced by atomic writes to the shared memory; while *PASSCoDe-Wild* proceeds without any locking and atomic operations, as a result of which the relationship (3) between primal and dual variables can be violated due to memory conflicts. Our contributions can be summarized below:

- We propose and analyze a family of asynchronous parallelization of the most efficient *DCD* algorithm: *PASSCoDe-Lock*, *PASSCoDe-Atomic*, *PASSCoDe-Wild*.
- We show linear convergence of *PASSCoDe-Atomic* under certain conditions.
- We present a *backward error analysis* for *PASSCoDe-Wild* and show that the converged solution is the exact solution of a primal problem with a perturbed regularizer. Therefore the performance is close-to-optimal on most of the datasets. To best of our knowledge, this is the first attempt to analyze a parallel machine learning algorithm with memory conflicts using backward error analysis, which is a standard tool in numerical analysis (Wilkinson, 1961).
- Experimental results show that our algorithms (*PASSCoDe-Atomic* and *PASSCoDe-Wild*) are much faster than existing methods. For example, on the `webspam` dataset, *PASSCoDe-Atomic* took 2 sec-

onds and *PASSCoDe-Wild* took 1.6 seconds to achieve 99% accuracy, while *CoCoA* took 11.5 seconds using 10 threads and *LIBLINEAR* took 10 seconds using 1 thread to achieve the same accuracy.

## 2 Related Work

**Stochastic Coordinate Descent.** Coordinate descent is a classical optimization technique that has been studied for a long time (Bertsekas, 1999; Luo & Tseng, 1992). Recently it has enjoyed renewed interest due to the success of “stochastic” coordinate descent in real applications (Hsieh et al., 2008; Nesterov, 2012). In terms of theoretical analysis, the convergence of (cyclic) coordinate descent has been studied for a long time (Luo & Tseng, 1992; Bertsekas, 1999), and the global linear convergence is presented recently under certain condition (Saha & Tewari, 2013; Wang & Lin, 2014).

**Stochastic Dual Coordinate Descent.** Many recent papers (Hsieh et al., 2008; Yu et al., 2011; Shalev-Shwartz & Zhang, 2013) have shown that solving the dual problem using coordinate descent algorithms is faster on large-scale datasets. The success of SDCD strongly relies on exploiting the primal-dual relationship (3) to speed up the gradient computation in the dual space. *DCD* has become the state-of-the-art solver implemented in *LIBLINEAR* (Fan et al., 2008). In terms of convergence of dual objective function, some standard theoretical guarantees for coordinate descent can be directly applied. Different from standard analysis, Shalev-Shwartz & Zhang (2013) presented the convergence rate in terms of duality gap.

**Parallel Stochastic Coordinate Descent.** In order to conduct coordinate updates in parallel, Richtárik & Takáč (2012) studied the algorithm where each processor updates a randomly selected block (or coordinate) simultaneously, and Bradley et al. (2011) proposed a similar algorithm for  $\ell_1$ -regularized problems. Scherrer et al. (2012) studied parallel greedy coordinate descent. However, the above synchronized methods usually face a trade-off in choosing the block size. If the block size is small, the load balancing problem leads to slow running time. If the block size is large, the convergence speed becomes much slower or the algorithm even diverges. These problems can be resolved by developing an asynchronous algorithm. Asynchronous coordinate descent has been studied by (Bertsekas & Tsitsiklis, 1989), but they require the Hessian to be diagonal dominant in order to establish the convergence. Recently, Liu et al. (2014); Liu & Wright (2014) proved linear convergence of asynchronous stochastic coordinate descent algorithms under the essential strong convexity condition and a “bounded staleness” condition, where they consider both “consistent read” and “inconsistent read” models. Avron et al. (2014) showed linear rate of convergence for the asynchronous randomized Gaussian-Seidel updates, which is a special case of coordinate descent on linear systems.

**Parallel Stochastic Dual Coordinate Descent.** For solving (5), each coordinate updates only requires the global primal variables  $\mathbf{w}$  and one local dual variable  $\alpha_i$ , thus algorithms only need to synchronize  $\mathbf{w}$ . Based on this observation, Yang (2013) proposed to update several coordinates or blocks simultaneously and update the global  $\mathbf{w}$ , and Jaggi et al. (2014) showed that each block can be solved with other approaches under the same framework. However, both these parallel *DCD* methods are synchronized algorithms.

To the best of our knowledge, this is the first to propose and analyze asynchronous parallel stochastic dual coordinate descent methods. By maintaining a primal solution  $\mathbf{w}$  while updating dual variables, our algorithm is much faster than the previous asynchronous coordinate descent methods of (Liu & Wright, 2014; Liu et al., 2014) for solving the dual problem (2). Our algorithms are also faster than synchronized dual coordinate descent methods (Yang, 2013; Jaggi et al., 2014) since the latest values of  $\mathbf{w}$  can be accessed by all the threads. In terms of theoretical contribution, the inconsistent read model in (Liu & Wright, 2014) cannot be directly applied to our algorithm because each update on  $\alpha_i$  is based on the shared  $\mathbf{w}$  vector. We further show linear convergence for *PASSCoDe-Atomic*, and study the properties of the converged solution for the

*wild* version of our algorithm (without any locking and atomic operations) using a backward error analysis. Our algorithm has been successfully applied to solve the collaborative ranking problem (Anonymous, 2015).

### 3 Algorithms

#### 3.1 Stochastic Dual Coordinate Descent

We first describe the Stochastic Dual Coordinate Descent (*DCD*) algorithm for solving the dual problem (2). At each iteration, *DCD* randomly picks a dual variable  $\alpha_i$  and updates it by minimizing the one variable subproblem (Eq. (4) in Algorithm 1). Without exploiting the structure of the quadratic term, the subproblems require substantial computation (need  $O(\text{nnz})$  time), where  $\text{nnz}$  is the total number of nonzero elements in the training data. However, if  $\mathbf{w}(\alpha)$  that satisfies (3) is maintained in memory, the subproblem  $D(\alpha + \delta \mathbf{e}_i)$  can be written as

$$D(\alpha + \delta \mathbf{e}_i) = \frac{1}{2} \|\mathbf{w} + \delta \mathbf{x}_i\|^2 + \ell_i^*(-(\alpha_i + \delta)),$$

and the optimal solution can be computed by

$$\delta = \arg \min_{\delta} \frac{1}{2} \left( \delta + \frac{\mathbf{w}^T \mathbf{x}_i}{\|\mathbf{x}_i\|^2} \right)^2 + \frac{1}{\|\mathbf{x}_i\|^2} \ell_i^*(-(\alpha_i + \delta)).$$

Note that all  $\|\mathbf{x}_i\|$  can be pre-computed and are constants. For each coordinate update we only need to solve a simple one-variable subproblem, and the main computation is in computing  $\mathbf{w}^T \mathbf{x}_i$ , which requires  $O(\text{nnz}/n)$  time. For SVM problems, the subproblem has a closed form solution, while for logistic regression problems it has to be solved by an iterative solver (see Yu et al. (2012) for details). The *DCD* algorithm, which is part of the popular *LIBLINEAR* package, is described in Algorithm 1.

---

#### Algorithm 1 Stochastic Dual Coordinate Descent (*DCD*)

---

**Input:** Initial  $\alpha$  and  $\mathbf{w} = \sum_{i=1}^n \alpha_i \mathbf{x}_i$

- 1: **while** not converged **do**
- 2:   Randomly pick  $i$
- 3:   Update  $\alpha_i \leftarrow \alpha_i + \Delta \alpha_i$ , where

$$\Delta \alpha_i \leftarrow \arg \min_{\delta} \frac{1}{2} \|\mathbf{w} + \delta \mathbf{x}_i\|^2 + \ell_i^*(-(\alpha_i + \delta)) \quad (4)$$

- 4:   Update  $\mathbf{w}$  by  $\mathbf{w} \leftarrow \mathbf{w} + \Delta \alpha_i \mathbf{x}_i$
  - 5: **end while**
- 

#### 3.2 Asynchronous Stochastic Dual Coordinate Descent

To parallelize *DCD* in a shared memory multi-core system, we propose a family of Asynchronous Stochastic Dual Coordinate Descent (*PASSCoDe*) algorithms. *PASSCoDe* is very simple but effective. Each thread repeatedly run the updates (steps 2 to 4) in Algorithm 1 using  $\mathbf{w}$ ,  $\alpha$ , and training data stored in a shared memory. The threads do not need to coordinate or synchronize their iterations. The details are shown in Algorithm 2.

---

**Algorithm 2** Parallel Asynchronous Stochastic dual Co-ordinate Descent (*PASSCoDe*)

---

**Input:** Initial  $\alpha$  and  $\mathbf{w} = \sum_{i=1}^n \alpha_i \mathbf{x}_i$

Each thread repeatedly performs the following updates:

step 1: Randomly pick  $i$

step 2: Update  $\alpha_i \leftarrow \alpha_i + \Delta\alpha_i$ , where

$$\Delta\alpha_i \leftarrow \arg \min_{\delta} \frac{1}{2} \|\mathbf{w} + \delta \mathbf{x}_i\|^2 + \ell_i^*(-(\alpha_i + \delta)) \quad (5)$$

step 3: Update  $\mathbf{w}$  by  $\mathbf{w} \leftarrow \mathbf{w} + \Delta\alpha_i \mathbf{x}_i$

---

Table 1: Scaling of *PASSCoDe* algorithms. We present the run time (in seconds) for each algorithm on the rcv1 dataset with 100 iterations, and the speedup of each method over the serial *DCD* algorithm (2x means it is two times faster than the serial algorithm).

Number of threads	Lock	Atomic	Wild
2	98.03s / 0.27x	15.28s / 1.75x	<b>14.08s / 1.90x</b>
4	106.11s / 0.25x	8.35s / 3.20x	<b>7.61s / 3.50x</b>
10	114.43s / 0.23x	3.86s / 6.91x	<b>3.59s / 7.43x</b>

Although *PASSCoDe* is a simple extension of *DCD* in a multi-core setting, there are many options in terms of locking/atomic operations for each step, and these choices lead to variations in speed and convergence properties, as we will show in this paper.

Note that the  $\Delta\alpha_i$  obtained by subproblem (5) is exactly the same as (4) in Algorithm 1 if only one thread is involved. However, when there are multiple threads, the  $\mathbf{w}$  vector may **not** be the latest one since some other threads have not completed the writes in step 3.

**PASSCoDe-Lock.** To ensure  $\mathbf{w} = \sum_i \alpha_i \mathbf{x}_i$  for the latest  $\alpha$ , we have to lock the following variables between step 1 and 2:

step 1.5: lock variables in  $N_i := \{w_t \mid (\mathbf{x}_i)_t \neq 0\}$ .

The locks are then released after step 3. With this locking mechanism, *PASSCoDe-Lock* will be serializable, i.e., generate the same solution sequence with the serial *DCD*. Unfortunately, threads will waste a lot of time due to the locks, so *PASSCoDe-Lock* is very slow comparing to the non-locking version (and even slower than the serial version of *DCD*). See Table 1 for details.

**PASSCoDe-Atomic.** The above locking scheme is to ensure that each thread updates  $\alpha_i$  based on the latest  $\mathbf{w}$  values. However, as shown in (Niu et al., 2011; Liu & Wright, 2014), the effect of using slightly stale values is usually limited in practice. Therefore, we propose an *PASSCoDe-Atomic* algorithm that avoids locking all the variables in  $N_i$  simultaneously. Instead, each thread just reads the current  $\mathbf{w}$  values from memory without any locking. In practice (see Section 5) we observe that the convergence speed is not significantly affected by using values of  $\mathbf{w}$ . However, to ensure that the limit point of the algorithm is still the global optimizer of (1), the equation  $\mathbf{w}^* = \sum_i \alpha_i^* \mathbf{x}_i$  has to be maintained. Therefore, we apply the following “atomic writes” in step 3:

step 3: For each  $j \in N(i)$

Update  $w_j \leftarrow w_j + \Delta\alpha_i (\mathbf{x}_i)_j$  **atomically**

Table 2: The performance of *PASSCoDe-Wild* using  $\hat{w}$  or  $\bar{w}$  for prediction. Results show that  $\hat{w}$  yields much better prediction accuracy, which justifies our theoretical analysis in Section 4.2.

	# threads	Prediction Accuracy (%) by		
		$\hat{w}$	$\bar{w}$	LIBLINEAR
news20	4	97.1	96.1	97.1
	8	97.2	93.3	
covtype	4	67.8	38.0	66.3
	8	67.6	38.0	
rcv1	4	97.7	97.5	97.7
	8	97.7	97.4	
webspam	4	99.1	93.1	99.1
	8	99.1	88.4	
kddb	4	88.8	79.7	88.8
	8	88.8	87.7	

*PASSCoDe-Atomic* is much faster than *PASSCoDe-Lock* as shown in Table 1 since the atomic writes for a single variable is much faster than locking all the variables. However, the convergence of *PASSCoDe-Atomic* is not guaranteed by any previous convergence analysis. To bridge this gap between practice and theory, we prove linear convergence of *PASSCoDe-Atomic* under certain conditions in Section 4.

**PASSCoDe-Wild.** Finally, we consider Algorithm 2 without any locks and atomic operations. The resulting algorithm, *PASSCoDe-Wild*, is faster than *PASSCoDe-Atomic* and *PASSCoDe-Lock* and can achieve almost linear speedup using a single processing unit. However, due to the memory conflicts in step 3, some of the "updates" to  $w$  will be over-written by other threads. As a result, the  $\hat{w}$  and  $\hat{\alpha}$  outputted by the algorithm usually do not satisfy Eq (3):

$$\hat{w} \neq \bar{w} := \sum_i \hat{\alpha}_i x_i, \quad (6)$$

where  $\hat{w}$ ,  $\hat{\alpha}$  are the primal and dual variables outputted by the algorithm, and  $\bar{w}$  defined in (6) is computed from  $\hat{\alpha}$ . It is easy to see that  $\hat{\alpha}$  is not the optimal solution of (2). Due to the same reason, in the prediction phase it is not clear whether  $\hat{w}$  or  $\bar{w}$  should be used. To answer this question, in Section 4 we show that  $\hat{w}$  is actually the optimal solution of a perturbed primal problem (1) using a backward error analysis, where the loss function is the same and the regularization term is slightly perturbed. As a result, the prediction should be done using  $\hat{w}$ , and this also yields much better performance in practice, as shown in Table 2 below.

We summarize the behavior of the three algorithms in Figure 1. Using locks, the algorithm *PASSCoDe-Lock* is serializable but very slow (even slower than the serial *DCD*). In the other extreme, the wild version without any lock and atomic operation has very good speed up, but the behavior can be totally different from the serial *DCD*. Luckily, in Section 4 we provide the convergence guarantee for *PASSCoDe-Atomic*, and apply a backward error analysis to show that *PASSCoDe-Wild* will converge to the solution with the same loss function with a slightly perturbed regularizer.

### 3.3 Implementation Details

**Deadlock Avoidance.** Without a proper implementation, the deadlock can arise in *PASSCoDe-Lock* because a thread needs to acquire all the locks associated with  $N_i$ . A simple way to avoid deadlock is by associating an ordering for all the locks such that each thread follows the same ordering to acquire the locks.



#### 4.1 Linear Convergence of PASSCoDe-Atomic

In PASSCoDe-Atomic, we assume all the updates before the  $(j - \tau)$ -th iteration has been written into  $\hat{\mathbf{w}}^j$ , therefore,

**Assumption 1.** *The set  $\mathcal{U}^j$  satisfies  $\mathcal{Z}^{j-\tau} \subseteq \mathcal{U}^j \subseteq \mathcal{Z}^j$ .*

Now we define some constants used in our theoretical analysis. Note that  $X \in \mathbb{R}^{n \times d}$  is the data matrix, and we use  $\bar{X} \in \mathbb{R}^{n \times d}$  to denote the normalized data matrix where each row is  $\bar{\mathbf{x}}_i^T = \mathbf{x}_i^T / \|\mathbf{x}_i\|^2$ . We then define

$$M_i = \max_{S \subseteq [d]} \left\| \sum_{t \in S} \bar{X}_{:,t} X_{i,t} \right\|, \quad M = \max_i M_i,$$

where  $[d] := \{1, \dots, d\}$  is the set of all the feature indices, and  $\bar{X}_{:,t}$  is the  $t$ -th column of  $\bar{X}$ . We also define  $L_{max}$  to be the Lipschitz constant of  $D(\cdot)$  within the level set  $\{\alpha \mid D(\alpha) \leq D(\alpha^0)\}$ ,  $R_{min} = \min_i \|\mathbf{x}_i\|^2$ ,  $R_{max} = \max_i \|\mathbf{x}_i\|^2$ . We assume that  $R_{max} = 1$  and there is no zero training sample, so  $R_{min} > 0$ .

To prove the convergence of asynchronous algorithms, we first show that the expected step size does not increase super-linearly by the following Lemma 1.

**Lemma 1.** *If  $\tau$  is small enough such that*

$$(6\tau(\tau + 1)^2 e M) / \sqrt{n} \leq 1, \quad (7)$$

*then PASSCoDe-Atomic satisfies the following inequality:*

$$E(\|\alpha^{j-1} - \alpha^j\|^2) \leq \rho E(\|\alpha^j - \alpha^{j+1}\|^2), \quad (8)$$

where  $\rho = (1 + \frac{6(\tau+1)eM}{\sqrt{n}})^2$ .

The detailed proof is in Appendix A.2. We use a similar technique as in (Liu & Wright, 2014) to prove this lemma, but the proof is different from (Liu & Wright, 2014) because

- Their “inconsistent read” model assumes  $\hat{\mathbf{w}}^j = \sum_i \dot{\alpha}_i \mathbf{x}_i$  for some  $\dot{\alpha}$ . However, in our case  $\hat{\mathbf{w}}^j$  may not be written in this form due to incomplete updates in step 3 of Algorithm 2.
- In (Liu & Wright, 2014), each coordinate is updated by  $\gamma \nabla_t f(\alpha)$  with a fixed step size  $\gamma$ . We consider the case that each subproblem (4) is solved exactly.

To show the linear convergence of our algorithms, we assume the objective function (2) satisfies the following property:

**Definition 1.** *The objective function (2) admits the global error bound if there is a constant  $\kappa$  such that*

$$\|\alpha - P_S(\alpha)\| \leq \kappa \|T(\alpha) - \alpha\|, \quad (9)$$

where  $P_S(\cdot)$  is the projection to the set of optimal solutions, and  $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$  is the operator defined by

$$T_t(\alpha) = \arg \min_u D(\alpha + (u - \alpha_t) \mathbf{e}_t) \quad \forall t = 1, \dots, n.$$

*The objective function satisfies the global error bound from the beginning if (9) holds for all  $\alpha$  satisfying*

$$D(\alpha) \leq D(\alpha^0)$$

*where  $\alpha^0$  is the initial point.*



This definition is a generalized version of Definition 6 in (Wang & Lin, 2014). We list several important machine learning problems that admit global error bounds:

- Support Vector Machines (SVM) with hinge loss (Boser et al., 1992):

$$\begin{aligned}\ell_i(z_i) &= C \max(1 - z_i, 0) \\ \ell_i^*(-\alpha_i) &= \begin{cases} -\alpha_i & \text{if } 0 \leq \alpha_i \leq C, \\ \infty & \text{otherwise.} \end{cases}\end{aligned}\tag{10}$$

- Support Vector Machines (SVM) with square hinge loss:

$$\begin{aligned}\ell_i(z_i) &= C \max(1 - z_i, 0)^2. \\ \ell_i^*(-\alpha_i) &= \begin{cases} -\alpha_i + \alpha_i^2/4C & \text{if } \alpha_i \geq 0, \\ \infty & \text{otherwise.} \end{cases}\end{aligned}\tag{11}$$

Note that  $C > 0$  is the penalty parameter that controls the weights between loss and regularization.

**Theorem 1.** *The Support Vector Machines (SVM) with hinge loss or square hinge loss satisfy the global error bound (9).*

*Proof.* For SVM with hinge loss, each element of the mapping  $T(\cdot)$  can be written as

$$\begin{aligned}T_t(\boldsymbol{\alpha}) &= \arg \min_u D(\boldsymbol{\alpha} + (u - \alpha_t)\mathbf{e}_t) \\ &= \arg \min_u \frac{1}{2} \|\mathbf{w}(\boldsymbol{\alpha}) + (u - \alpha_t)\mathbf{x}_t\|^2 + \ell^*(-u) \\ &= \Pi_{\mathbb{X}}\left(\frac{\mathbf{w}(\boldsymbol{\alpha})^T \mathbf{x}_t - 1}{\|\mathbf{x}_t\|^2}\right) = \Pi_{\mathbb{X}}\left(\frac{\nabla_t D(\boldsymbol{\alpha})}{\|\mathbf{x}_t\|^2}\right),\end{aligned}$$

where  $\Pi_{\mathbb{X}}$  is the projection to the set  $\mathbb{X}$ , and for hinge-loss SVM  $\mathbb{X} := [0, C]$ . Using Lemma 26 in (Wang & Lin, 2014), we can show that for all  $t = 1, \dots, n$

$$\begin{aligned}|\alpha_t - \Pi_{\mathbb{X}}\left(\frac{\nabla_t D(\boldsymbol{\alpha})}{\|\mathbf{x}_t\|^2}\right)| &\geq \min(1, \frac{1}{\|\mathbf{x}_t\|^2}) |\alpha_t - \Pi_{\mathbb{X}}(\nabla_t D(\boldsymbol{\alpha}))| \\ &\geq \min(1, \frac{1}{R_{max}^2}) |\alpha_t - \Pi_{\mathbb{X}}(\nabla_t D(\boldsymbol{\alpha}))| \\ &\geq |\alpha_t - \Pi_{\mathbb{X}}(\nabla_t D(\boldsymbol{\alpha}))|,\end{aligned}$$

Table 3: Data statistics.  $\tilde{n}$  is the number of test instances.  $\bar{d}$  is the average nnz per instance.

	$n$	$\tilde{n}$	$d$	$\bar{d}$	$C$
news20	16,000	3,996	1,355,191	455.5	2
covtype	500,000	81,012	54	11.9	0.0625
rcv1	677,399	20,242	47,236	73.2	1
webspam	280,000	70,000	16,609,143	3727.7	1
kdddb	19,264,097	748,401	29,890,095	29.4	1

where the last inequality is due to the assumption that  $R_{max} = 1$ . Therefore,

$$\begin{aligned}
\|\alpha - T(\alpha)\|_2 &\geq \frac{1}{\sqrt{n}} \|\alpha - T(\alpha)\|_1 \\
&\geq \frac{1}{\sqrt{n}} \sum_{t=1}^n |\alpha_t - \Pi_{\mathbb{X}}(\nabla_t D(\alpha))| \\
&= \frac{1}{\sqrt{n}} \|\nabla^+ D(\alpha)\|_1 \\
&\geq \frac{1}{\sqrt{n}} \|\nabla^+ D(\alpha)\|_2 \\
&\geq \frac{1}{\kappa_0 \sqrt{n}} \|\alpha - P_S(\alpha)\|_2,
\end{aligned}$$

where  $\nabla^+ D(\alpha)$  is the projected gradient defined in Definition 5 of (Wang & Lin, 2014) and  $\kappa_0$  is the  $\kappa$  defined in Theorem 18 of (Wang & Lin, 2014). Thus, with  $\kappa = \kappa_0 \sqrt{n}$ , we obtain that the dual function of the hinge-loss SVM satisfies the global error bound defined in Definition 1. Similarly, we can show that the SVM with squared-hinge loss satisfies the global error bound.  $\square$

Next we explicitly state the linear convergence guarantee for *PASSCoDe-Atomic*.

**Theorem 2.** *Assume the objective function (2) admits a global error bound from the beginning and the Lipschitz constant  $L_{max}$  is finite in the level set. If (7) holds and*

$$1 \geq \frac{2L_{max}}{R_{min}^2} \left(1 + \frac{e\tau M}{\sqrt{n}}\right) \left(\frac{\tau^2 M^2 e^2}{n}\right)$$

*then PASSCoDe-Atomic has a global linear convergence rate in expectation, that is,*

$$E[D(\alpha^{j+1})] - D(\alpha^*) \leq \eta (E[D(\alpha^j)] - D(\alpha^*)), \quad (12)$$

*where  $\alpha^*$  is the optimal solution and*

$$\eta = 1 - \frac{\kappa}{L_{max}} \left(1 - \frac{2L_{max}}{R_{min}^2} \left(1 + \frac{e\tau M}{\sqrt{n}}\right) \left(\frac{\tau^2 M^2 e^2}{n}\right)\right) \quad (13)$$

## 4.2 Backward Error Analysis for PASSCoDe-Wild

In *PASSCoDe-Wild*, assume the sequence  $\{\alpha^j\}$  converges to  $\hat{\alpha}$  and  $\{\mathbf{w}^j\}$  converges to  $\hat{\mathbf{w}}$ . Now we show that the dual solution  $\hat{\alpha}$  and the corresponding primal variables  $\bar{\mathbf{w}} = \sum_{i=1}^n \hat{\alpha}_i \mathbf{x}_i$  are actually the dual and primal solutions of a perturbed problem:

**Theorem 3.**  $\hat{\alpha}$  is the optimal solution of a perturbed dual problem

$$\hat{\alpha} = \arg \min_{\alpha} D(\alpha) - \sum_{i=1}^n \alpha_i \epsilon^T \mathbf{x}_i, \quad (14)$$

and  $\bar{\mathbf{w}} = \sum_i \hat{\alpha}_i \mathbf{x}_i$  is the solution of the corresponding primal problem:

$$\bar{\mathbf{w}} = \arg \min_{\mathbf{w}} \frac{1}{2} \mathbf{w}^T \mathbf{w} + \sum_{i=1}^n \ell_i((\mathbf{w} - \epsilon)^T \mathbf{x}_i), \quad (15)$$

where  $\epsilon \in \mathbb{R}^d$  is given by  $\epsilon = \bar{\mathbf{w}} - \hat{\mathbf{w}}$ .

*Proof.* By definition,  $\hat{\alpha}$  is the limit point of *PASSCoDe-Wild*. Therefore,  $\{\Delta \alpha_i\} \rightarrow 0$  for all  $i$ . Combining with the fact that  $\{\hat{\mathbf{w}}^j\} \rightarrow \hat{\mathbf{w}}$ , we have

$$-\hat{\mathbf{w}}^T \mathbf{x}_i \in \partial_{\alpha_i} \ell_i^*(-\hat{\alpha}_i), \quad \forall i.$$

Since  $\hat{\mathbf{w}} = \bar{\mathbf{w}} - \epsilon$ , we have

$$\begin{aligned} -(\bar{\mathbf{w}} - \epsilon)^T \mathbf{x}_i &\in \partial_{\alpha_i} \ell_i^*(-\hat{\alpha}_i), \quad \forall i \\ -\bar{\mathbf{w}}^T \mathbf{x}_i &\in \partial_{\alpha_i} (\ell_i^*(-\hat{\alpha}_i) - \hat{\alpha}_i \epsilon^T \mathbf{x}_i), \quad \forall i \\ 0 &\in \partial_{\alpha_i} \left( \frac{1}{2} \left\| \sum_{i=1}^n \hat{\alpha}_i \mathbf{x}_i \right\|^2 + \ell_i^*(-\hat{\alpha}_i) - \hat{\alpha}_i \epsilon^T \mathbf{x}_i \right), \quad \forall i \end{aligned}$$

which is the optimality condition of (14). Thus,  $\hat{\alpha}$  is the optimal solution of (14).

For the second part of the theorem, let's consider the following equivalent primal problem and its Lagrangian:

$$\begin{aligned} \min_{\mathbf{w}, \xi} \quad & \frac{1}{2} \mathbf{w}^T \mathbf{w} + \sum_{i=1}^n \ell_i(\xi_i) \quad \text{s.t. } \xi_i = (\mathbf{w} - \epsilon)^T \mathbf{x}_i \quad \forall i = 1, \dots, n \\ L(\mathbf{w}, \xi, \alpha) := \quad & \frac{1}{2} \mathbf{w}^T \mathbf{w} + \sum_{i=1}^n \{ \ell_i(\xi_i) + \alpha_i (\xi_i - \mathbf{w}^T \mathbf{x}_i + \epsilon^T \mathbf{x}_i) \} \end{aligned}$$

The corresponding convex version of the dual function can be derived as follows.

$$\begin{aligned} \hat{D}(\alpha) &= \max_{\mathbf{w}, \xi} -L(\mathbf{w}, \xi, \alpha) \\ &= \left( \max_{\mathbf{w}} -\frac{1}{2} \mathbf{w}^T \mathbf{w} + \sum_{i=1}^n \alpha_i \mathbf{w}^T \mathbf{x}_i \right) + \sum_{i=1}^n \left( \max_{\xi_i} -\ell_i(\xi_i) - \alpha_i \xi_i \right) - \alpha_i \epsilon^T \mathbf{x}_i \\ &= \frac{1}{2} \left\| \sum_{i=1}^n \alpha_i \mathbf{x}_i \right\|^2 + \sum_{i=1}^n \ell_i^*(-\alpha_i) - \alpha_i \epsilon^T \mathbf{x}_i \\ &= D(\alpha) - \sum_{i=1}^n \alpha_i \epsilon^T \mathbf{x}_i \end{aligned}$$

The last second equality comes from 1) the substitution of  $\mathbf{w}^* = \sum_{i=1}^T \alpha_i \mathbf{x}_i$  obtained by setting  $\nabla_{\mathbf{w}} - L(\mathbf{w}, \xi, \alpha) = 0$ ; 2) the definition of the conjugate function  $\ell_i^*(-\alpha_i)$ . Thus, the second part of the theorem follows.  $\square$

Note that  $\epsilon$  is the error caused by the memory conflicts. From Theorem 3,  $\bar{\mathbf{w}}$  is the optimal solution of the “biased” primal problem (15), however, in (15) the actual model that fits the loss function should be  $\hat{\mathbf{w}} = \bar{\mathbf{w}} - \epsilon$ . Therefore after the training process we should use  $\hat{\mathbf{w}}$  to predict, which is the  $\mathbf{w}$  we maintained during the parallel coordinate descent updates. Replacing  $\mathbf{w}$  by  $\mathbf{w} - \epsilon$  in (15), we have the following corollary :

**Corollary 1.**  *$\hat{\mathbf{w}}$  computed by PASSCoDe-Wild is the solution of the following perturbed primal problem:*

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w}} \frac{1}{2}(\mathbf{w} + \epsilon)^T(\mathbf{w} + \epsilon) + \sum_{i=1}^n \ell_i(\mathbf{w}^T \mathbf{x}_i) \quad (16)$$

The above corollary shows that the computed primal solution  $\hat{\mathbf{w}}$  is actually the *exact* solution of a perturbed problem (where the perturbation is on the regularizer). This strategy (of showing that the computed solution to a problem is the exact solution of a perturbed problem) is inspired by the *backward error analysis* technique commonly employed in numerical analysis (Wilkinson, 1961)<sup>1</sup>.

## 5 Experimental Results

We conduct several experiments and show that the proposed *PASSCoDe-Atomic* and *PASSCoDe-Wild* have superior performance compared to other state-of-the-art parallel coordinate descent algorithms. We consider the hinge loss and five datasets: `news20`, `covtype`, `rcv1`, `webspam`, and `kdddb`. Detailed information is shown in Table 3. To have a fair comparison, we implement all compared methods in C++ using OpenMP as the parallel programming framework. All the experiments are performed on an Intel multi-core dual-socket machine with 256 GB memory. Each socket is associated with 10 computation cores. We explicitly enforce that all the threads use cores from the same socket to avoid inter-socket communication. Our codes will be publicly available. We focus on solving the (hinge loss) SVM (see (5) in the Appendix) in the experiments, but the algorithms can also be applied to other objective functions. Note that some of the figures are in Appendix 6.

### Serial Baselines.

- *DCD*: we implement Algorithm 1. Instead of sampling with replacement, a random permutation is used to enforce random sampling without replacement.
- *LIBLINEAR*: we use the implementation in <http://www.csie.ntu.edu.tw/~cjlin/liblinear>. This implementation is equivalent to *DCD* with the shrinking strategy.

### Compared Parallel Implementation.

- *PASSCoDe*: We implement the proposed three variants of Algorithm 2 using *DCD* as the building block: *Wild*, *Atomic*, and *Lock*.
- *CoCoA*: We implement a multi-core version of *CoCoA* (Jaggi et al., 2014) with  $\beta_K = 1$  and *DCD* as its local dual method.
- *AsySCD*: We follow the description in (Liu & Wright, 2014; Liu et al., 2014) to implement *AsySCD* with the step length  $\gamma = \frac{1}{2}$  and the shuffling period  $p = 10$  as suggested in (Liu et al., 2014).

---

<sup>1</sup>J. H. Wilkinson received the Turing Award in 1970, partly for his work on backward error analysis

## 5.1 Convergence in terms of iterations.

The primal objective function value is used to determine the convergence. Note that we still use  $P(\hat{w})$  for *PASSCoDe-Wild*, although the true primal objective should be (16). As long as  $\hat{w}^T \epsilon$  remains small enough, the trend of (16) and  $P(\hat{w})$  are similar.

Figure 4(a), 5(a), 6(a) show the convergence results of *PASSCoDe-Wild*, *PASSCoDe-Atomic*, *CoCoA*, and *AsySCD* with 10 threads in terms of number of iterations. The horizontal line in grey indicates the primal objective function value obtained by LIBLINEAR using the default stopping condition. The result for LIBLINEAR is also included for reference. We have the follow observations

- Convergence of three *PASSCoDe* variants are almost identical and very close to the convergence behavior of serial LIBLINEAR on three large sparse datasets (rcv1, webspam, and kddb).
- *PASSCoDe-Wild* and *PASSCoDe-Atomic* converge significantly faster than *CoCoA*.
- On *covtype*, a more dense dataset, all three algorithms (*PASSCoDe-Wild*, *PASSCoDe-Atomic*, and *CoCoA*) have slower convergence.

## 5.2 Efficiency.

**Timing.** To have a fair comparison, we include both initialization and computation into the timing results. For *DCD*, *PASSCoDe*, *CoCoA*, initialization takes one pass of entire data matrix (which is  $O(nnz(X))$ ) to compute  $\|x_i\|$  for each instance. In the initialization stage, *AsySCD* requires  $O(n \times nnz(X))$  time and  $O(n^2)$  space to form and store the Hessian matrix  $Q$  for (2). Thus, we only have results on *news20* for *AsySCD* as all other datasets are too large for *AsySCD* to fit  $Q$  in even 256 GB memory. Note that we also parallelize the initialization part for each algorithm in our implementation to have a fair comparison.

Figures 2(b), 3(b), 4(b), 5(b), 6(b) show the primal objective values in terms of time and Figures 2(c), 3(c), 4(c), 5(c), 6(c) shows the accuracy in terms of time. Note that the x-axis for *news20*, *covtype*, and *rcv1* is in log-scale. A horizontal line in gray in each figure denotes the objective values/accuracy obtained by LIBLINEAR using the default stopping condition. We have the following observations:

- From Figures 4(b) and 4(c), we can see that *AsySCD* is orders of magnitude slower than other approaches including parallel methods and serial reference (*AsySCD* using 10 cores takes 0.4 seconds to run 10 iterations, while all the other parallel approaches takes less than 0.14 seconds, and LIBLINEAR takes less than 0.3 seconds). In fact, *AsySCD* is still slower than other methods even when the initialization time is excluded. This is expected because *AsySCD* is a parallel version of a standard coordinate descent method, which is known to be much slower than *DCD* for (2). Since *AsySCD* runs out of memory for all the other larger datasets, we do not show the results in other figures.
- In most figures, both *PASSCoDe* approaches outperform *CoCoA*. In Figure 6(c), *kddb* shows better accuracy performance in the early stage which can be explained by the ensemble nature of *CoCoA*. In the long term, it still converges to the accuracy obtained by LIBLINEAR.
- For all datasets, *PASSCoDe-Wild* is shown to be slightly faster than *PASSCoDe-Atomic*. Given the fact that both methods show similar convergence in terms of iterations, this phenomenon can be explained by the effect of atomic operations. We can observe that more dense the dataset, larger the difference between *PASSCoDe-Wild* and *PASSCoDe-Atomic*.

## 5.3 Speedup

We are interested in the following evaluation criterion:

$$\text{speedup} := \frac{\text{time taken by the target method with } p \text{ threads}}{\text{time taken by the best serial reference method}},$$

This criterion is different from *scaling*, where the denominator is replaced by “time taken for the target method with single thread.” Note that a method can have perfect scaling but very poor speedup. Figures 2(d), 3(d), 4(d), 5(d), 6(d) shows the speedup results, where 1) *DCD* is used as the best serial reference; 2) the shrinking heuristic is turned off for all *PASSCoDe* and *DCD* to have fair comparison; 3) the initialization time is excluded from the computation of speedup.

- *PASSCoDe-Wild* has very good speedup performance compared to other approaches. It achieves about 6 to 8 speedups using 10 threads on all the datasets.
- From Figure 2(d), we can see that *AsySCD* does not have any “speedup” over the serial reference, although it is shown to have almost linear scaling (Liu et al., 2014; Liu & Wright, 2014).

## 6 Conclusions

In this paper, we present a family of parallel asynchronous stochastic dual coordinate descent algorithms in the shared memory multi-core setting, where each thread repeatedly selects a random dual variable and conducts coordinate updates using the primal variables that are stored in the shared memory. We analyze the convergence properties when different locking/atomic mechanism is used. For the setting with atomic updates, we show the linear convergence under certain condition. For the setting without any lock or atomic write, which achieves the best speed up, we present a backward error analysis to show that the primal variables obtained by the algorithm is the exact solution for a primal problem with perturbed regularizer. Experimental results show that our algorithms are much faster than previous parallel coordinate descent solvers.





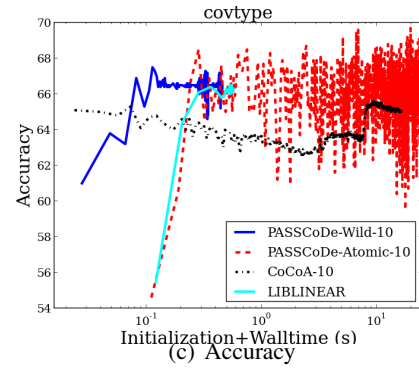









Figure 2: news20 dataset

Figure 3: covtype dataset

Figure 4: rcv1 dataset






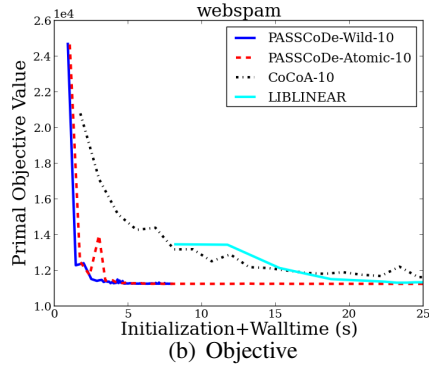




Figure 5: webspam dataset

Figure 6: kddb dataset



## References

- Anonymous. Preference completion: Large-scale collaborative ranking from pairwise comparison. submitted to ICML, 2015.
- Avron, H., Druinsky, A., and Gupta, A. Revisiting asynchronous linear solvers: Provable convergence rate through randomization. In *IEEE International Parallel and Distributed Processing Symposium*, 2014.
- Bertsekas, Dimitri P. *Nonlinear Programming*. Athena Scientific, Belmont, MA 02178-9998, second edition, 1999.
- Bertsekas, Dimitri P. and Tsitsiklis, John N. *Parallel and Distributed Computation: Numerical Methods*. Prentice Hall, 1989.
- Boser, Bernhard E., Guyon, Isabelle, and Vapnik, Vladimir. A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, pp. 144–152. ACM Press, 1992.
- Bradley, Joseph K., Kyrola, Aapo, Bickson, Danny, and Guestrin, Carlos. Parallel coordinate descent for  $l_1$ -regularized loss minimization. In *ICML*, 2011.
- Chang, Kai-Wei, Hsieh, Cho-Jui, and Lin, Chih-Jen. Coordinate descent method for large-scale L2-loss linear SVM. *Journal of Machine Learning Research*, 9:1369–1398, 2008. URL <http://www.csie.ntu.edu.tw/~cjlin/papers/cdl2.pdf>.
- Fan, Rong-En, Chang, Kai-Wei, Hsieh, Cho-Jui, Wang, Xiang-Rui, and Lin, Chih-Jen. LIBLINEAR: a library for large linear classification. *JMLR*, 9:1871–1874, 2008.
- Hsieh, Cho-Jui, Chang, Kai-Wei, Lin, Chih-Jen, Keerthi, S. Sathiya, and Sundararajan, Sellamanickam. A dual coordinate descent method for large-scale linear SVM. In *Proceedings of the Twenty Fifth International Conference on Machine Learning (ICML)*, 2008. URL <http://www.csie.ntu.edu.tw/~cjlin/papers/cddual.pdf>.
- Huang, Fang-Lan, Hsieh, Cho-Jui, Chang, Kai-Wei, and Lin, Chih-Jen. Iterative scaling and coordinate descent methods for maximum entropy. In *Proceedings of the 47th Annual Meeting of the Association of Computational Linguistics (ACL)*, 2009. Short paper.
- Jaggi, Martin, Smith, Virginia, Takáč, Martin, Terhorst, Jonathan, Hofmann, Thomas, and Jordan, Michael I. Communication-efficient distributed dual coordinate ascent. In *Advances in Neural Information Processing Systems* 27. 2014.
- Keerthi, S. Sathiya, Sundararajan, Sellamanickam, Chang, Kai-Wei, Hsieh, Cho-Jui, and Lin, Chih-Jen. A sequential dual method for large scale multi-class linear SVMs. In *Proceedings of the Forteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 408–416, 2008. URL [http://www.csie.ntu.edu.tw/~cjlin/papers/sdm\\_kdd.pdf](http://www.csie.ntu.edu.tw/~cjlin/papers/sdm_kdd.pdf).
- Lin, Chih-Jen, Weng, Ruby C., and Keerthi, S. Sathiya. Trust region Newton method for large-scale logistic regression. In *Proceedings of the 24th International Conference on Machine Learning (ICML)*, 2007. Software available at <http://www.csie.ntu.edu.tw/~cjlin/liblinear>.

- Liu, J. and Wright, S. J. Asynchronous stochastic coordinate descent: Parallelism and convergence properties. 2014. URL <http://arxiv.org/abs/1403.3862>.
- Liu, J., Wright, S. J., Re, C., and Bittorf, V. An asynchronous parallel stochastic coordinate descent algorithm. In *ICML*, 2014.
- Luo, Zhi-Quan and Tseng, Paul. On the convergence of coordinate descent method for convex differentiable minimization. *Journal of Optimization Theory and Applications*, 72(1):7–35, 1992.
- Nesterov, Yurii E. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- Niu, Feng, Recht, Benjamin, Ré, Christopher, and Wright, Stephen J. HOGWILD!: a lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems 24*, pp. 693–701, 2011.
- Richtárik, Peter and Takáč, Martin. Parallel coordinate descent methods for big data optimization. *Mathematical Programming*, 2012. Under revision.
- Saha, Ankan and Tewari, Ambuj. On the nonasymptotic convergence of cyclic coordinate descent methods. *SIAM Journal on Optimization*, 23(1):576–601, 2013.
- Scherrer, C., Tewari, A., Halappanavar, M., and Haglin, D. Feature clustering for accelerating parallel coordinate descent. In *NIPS*, 2012.
- Shalev-Shwartz, S., Singer, Y., and Srebro, N. Pegasos: primal estimated sub-gradient solver for SVM. In *ICML*, 2007.
- Shalev-Shwartz, Shai and Zhang, Tong. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14:567–599, 2013.
- Wang, Po-Wei and Lin, Chih-Jen. Iteration complexity of feasible descent methods for convex optimization. *Journal of Machine Learning Research*, 15:1523–1548, 2014. URL <http://www.csie.ntu.edu.tw/~cjlin/papers/cdlinear.pdf>.
- Wilkinson, J. H. Error analysis of direct methods of matrix inversion. *Journal of the ACM*, 1961.
- Yang, T. Trading computation for communication: Distributed stochastic dual coordinate ascent. In *NIPS*, 2013.
- Yu, Hsiang-Fu, Huang, Fang-Lan, and Lin, Chih-Jen. Dual coordinate descent methods for logistic regression and maximum entropy models. *Machine Learning*, 85(1-2):41–75, October 2011. URL [http://www.csie.ntu.edu.tw/~cjlin/papers/maxent\\_dual.pdf](http://www.csie.ntu.edu.tw/~cjlin/papers/maxent_dual.pdf).
- Yu, Hsiang-Fu, Hsieh, Cho-Jui, Si, Si, and Dhillon, Inderjit S. Scalable coordinate descent approaches to parallel matrix factorization for recommender systems. In *Proceedings of the IEEE International Conference on Data Mining*, pp. 765–774, 2012.
- Zhang, Tong. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *Proceedings of the 21th International Conference on Machine Learning (ICML)*, 2004.

## A Linear Convergence for *PASSCoDe-Atomic*

### A.1 Notations and Prepositions

#### A.1.1 Notations

- For all  $i = 1, \dots, n$ , we have the following definitions:

$$\begin{aligned} h_i(u) &:= \frac{\ell_i^*(-u)}{\|\mathbf{x}_i\|^2} \\ \text{prox}_i(s) &:= \arg \min_u \frac{1}{2}(u - s)^2 + h_i(u) \\ T_i(\mathbf{w}, s) &:= \arg \min_u \frac{1}{2}\|\mathbf{w} + (u - s)\mathbf{x}_i\|^2 + \ell_i^*(-u) \\ &= \arg \min_u \frac{1}{2} \left[ u - \left( s - \frac{\mathbf{w}^T \mathbf{x}_i}{\|\mathbf{x}_i\|^2} \right) \right]^2 + h_i(u), \end{aligned}$$

where  $\mathbf{w} \in R^d$  and  $s \in R$ . We also denote  $\text{prox}(s)$  as the proximal operator from  $R^n$  to  $R^n$  such that  $(\text{prox}(\mathbf{x}))_i = \text{prox}_i(s_i)$ . We can see the connection of the above operator and the proximal operator:  $T_i(\mathbf{w}, s) = \text{prox}_i(s - \frac{\mathbf{w}^T \mathbf{x}_i}{\|\mathbf{x}_i\|^2})$ .

- Let  $\{\boldsymbol{\alpha}^j\}$  and  $\{\hat{\mathbf{w}}^j\}$  be the sequence generated/maintained by Algorithm 2 using

$$\alpha_t^{j+1} = \begin{cases} T_t(\hat{\mathbf{w}}^j, \alpha_t^j) & \text{if } t = i(j), \\ \alpha_t^j & \text{if } t \neq i(j), \end{cases}$$

where  $i(j)$  is the index selected at  $j$ -th iteration. For convenience, we define

$$\Delta \alpha_j = \alpha_{i(j)}^{j+1} - \alpha_{i(j)}^j.$$

- Let  $\{\tilde{\boldsymbol{\alpha}}^j\}$  be the sequence defined by

$$\tilde{\alpha}_t^{j+1} = T_t(\hat{\mathbf{w}}^j, \alpha_t^j) \quad \forall t = 1, \dots, n.$$

Note that  $\tilde{\alpha}_{i(j)}^{j+1} = \alpha_{i(j)}^{j+1}$  and  $\tilde{\boldsymbol{\alpha}}^{j+1} = \text{prox}(\boldsymbol{\alpha}^j - \bar{X} \hat{\mathbf{w}}^j)$ .

- Let  $\bar{\mathbf{w}}^j = \sum_i \alpha_i^j \mathbf{x}_i$  be the “true” primal variables corresponding to  $\boldsymbol{\alpha}^j$ .

#### A.1.2 Prepositions

##### Proposition 1.

$$E_{i(j)}(\|\boldsymbol{\alpha}^{j+1} - \boldsymbol{\alpha}^j\|^2) = \frac{1}{n} \|\tilde{\boldsymbol{\alpha}}^{j+1} - \boldsymbol{\alpha}^j\|^2. \quad (17)$$

*Proof.* It can be proved by the definition of  $\tilde{\boldsymbol{\alpha}}$  and the assumption that  $i(j)$  is uniformly random selected from  $\{1, \dots, n\}$ .  $\square$

**Preposition 2.**

$$\|\bar{X}\bar{\mathbf{w}}^j - \bar{X}\hat{\mathbf{w}}^j\| \leq M \sum_{t=j-\tau}^{j-1} |\Delta\alpha_t|. \quad (18)$$

*Proof.*

$$\begin{aligned} \|\bar{X}\bar{\mathbf{w}}^j - \bar{X}\hat{\mathbf{w}}^j\| &= \|\bar{X}(\sum_{(t,k) \in \mathcal{Z}^j \setminus \mathcal{U}^j} (\Delta\alpha_t) X_{i(t),k} \mathbf{e}_k)\| = \|\sum_{(t,k) \in \mathcal{Z}^j \setminus \mathcal{U}^j} (\Delta\alpha_t) \bar{X}_{:,k} X_{i(t),k}\| \\ &\leq \sum_{t=j-1}^{j-\tau} |\Delta\alpha_t| M_i \leq M \sum_{t=j-\tau}^{j-1} |\Delta\alpha_t| \end{aligned}$$

□

**Preposition 3.** For any  $\mathbf{w}_1, \mathbf{w}_2 \in R^d$  and  $s_1, s_2 \in R$ ,

$$|T_i(\mathbf{w}_1, s_1) - T_i(\mathbf{w}_2, s_2)| \leq |s_1 - s_2 + \frac{(\mathbf{w}_1 - \mathbf{w}_2)^T \mathbf{x}_i}{\|\mathbf{x}_i\|^2}|. \quad (19)$$

*Proof.* It can be proved by the connection of  $T_i(\mathbf{w}, s)$  and  $\text{prox}_i(\cdot)$  and the non-expansiveness of the proximal operator. □

**Preposition 4.** Let  $M \geq 1$ ,  $q = \frac{6(\tau+1)eM}{\sqrt{n}}$ ,  $\rho = (1+q)^2$ , and  $\theta = \sum_{t=1}^{\tau} \rho^{t/2}$ . If  $M \geq 1$  and  $q(\tau+1) \leq 1$ , then  $\rho^{(\tau+1)/2} \leq e$ , and

$$\rho^{-1} \leq 1 - \frac{4 + 4M + 4M\theta}{\sqrt{n}}. \quad (20)$$

*Proof.* By the definition of  $\rho$  and the condition  $q(\tau+1) \leq 1$ , we have

$$\rho^{(\tau+1)/2} = \left( \left( \rho^{1/2} \right)^{1/q} \right)^{q(\tau+1)} = \left( (1+q)^{1/q} \right)^{q(\tau+1)} \leq e^{q(\tau+1)} \leq e.$$

By the definitions of  $q$ , we know that

$$q = \rho^{1/2} - 1 = \frac{6(\tau+1)eM}{\sqrt{n}} \Rightarrow \frac{3}{2} = \frac{\sqrt{n}(\rho^{1/2} - 1)}{4(\tau+1)eM}.$$

We can derive

$$\begin{aligned} \frac{3}{2} &= \frac{\sqrt{n}(\rho^{1/2} - 1)}{4(\tau+1)eM} \\ &\leq \frac{\sqrt{n}(\rho^{1/2} - 1)}{4(\tau+1)\rho^{(\tau+1)/2}M} && \because \rho^{(\tau+1)/2} \leq e \\ &\leq \frac{\sqrt{n}(\rho^{1/2} - 1)}{4(1+\theta)\rho^{1/2}M} && \because 1 + \theta = \sum_{t=0}^{\tau} \rho^{t/2} \leq (\tau+1)\rho^{\tau/2} \\ &= \frac{\sqrt{n}(1 - \rho^{-1/2})}{4(1+\theta)M} \\ &\leq \frac{\sqrt{n}(1 - \rho^{-1})}{4(1+\theta)M} && \because \rho^{-1/2} \leq 1 \end{aligned}$$

Combining the condition that  $M \geq 1$  and  $1 + \theta \geq 1$ , we have

$$\frac{\sqrt{n}(1 - \rho^{-1}) - 4}{4(1 + \theta)M} \geq \frac{\sqrt{n}(1 - \rho^{-1})}{4(1 + \theta)M} - \frac{1}{2} \geq 1,$$

which leads to

$$\begin{aligned} 4(1 + \theta)M &\leq \sqrt{n} - \sqrt{n}\rho^{-1} - 4 \\ \rho^{-1} &\leq 1 - \frac{4 + 4M + 4M\theta}{\sqrt{n}}. \end{aligned}$$

□

## A.2 Proof of Lemma 1

Similar to Liu & Wright (2014), we prove Eq. (8) by induction. First, we know that for any two vectors  $\mathbf{a}$  and  $\mathbf{b}$ , we have

$$\|\mathbf{a}\|^2 - \|\mathbf{b}\|^2 \leq 2\|\mathbf{a}\|\|\mathbf{b} - \mathbf{a}\|.$$

See Liu & Wright (2014) for a proof for the above inequality. Thus, for all  $j$ , we have

$$\|\boldsymbol{\alpha}^{j-1} - \tilde{\boldsymbol{\alpha}}^j\|^2 - \|\boldsymbol{\alpha}^j - \tilde{\boldsymbol{\alpha}}^{j+1}\|^2 \leq 2\|\boldsymbol{\alpha}^{j-1} - \tilde{\boldsymbol{\alpha}}^j\|\|\boldsymbol{\alpha}^j - \tilde{\boldsymbol{\alpha}}^{j+1} - \boldsymbol{\alpha}^{j-1} + \tilde{\boldsymbol{\alpha}}^j\|. \quad (21)$$

The second factor in the r.h.s of (21) is bounded as follows:

$$\begin{aligned} &\|\boldsymbol{\alpha}^j - \tilde{\boldsymbol{\alpha}}^{j+1} - \boldsymbol{\alpha}^{j-1} + \tilde{\boldsymbol{\alpha}}^j\| \\ &\leq \|\boldsymbol{\alpha}^j - \boldsymbol{\alpha}^{j-1}\| + \|\text{prox}(\boldsymbol{\alpha}^j - \bar{X}\hat{\mathbf{w}}^j) - \text{prox}(\boldsymbol{\alpha}^{j-1} - \bar{X}\hat{\mathbf{w}}^{j-1})\| \\ &\leq \|\boldsymbol{\alpha}^j - \boldsymbol{\alpha}^{j-1}\| + \|(\boldsymbol{\alpha}^j - \bar{X}\hat{\mathbf{w}}^j) - (\boldsymbol{\alpha}^{j-1} - \bar{X}\hat{\mathbf{w}}^{j-1})\| \\ &\leq \|\boldsymbol{\alpha}^j - \boldsymbol{\alpha}^{j-1}\| + \|\boldsymbol{\alpha}^j - \boldsymbol{\alpha}^{j-1}\| + \|\bar{X}\hat{\mathbf{w}}^j - \bar{X}\hat{\mathbf{w}}^{j-1}\| \\ &= 2\|\boldsymbol{\alpha}^j - \boldsymbol{\alpha}^{j-1}\| + \|\bar{X}\hat{\mathbf{w}}^j - \bar{X}\hat{\mathbf{w}}^{j-1}\| \\ &= 2\|\boldsymbol{\alpha}^j - \boldsymbol{\alpha}^{j-1}\| + \|\bar{X}\hat{\mathbf{w}}^j - \bar{X}\bar{\mathbf{w}}^j + \bar{X}\bar{\mathbf{w}}^j - \bar{X}\bar{\mathbf{w}}^{j-1} + \bar{X}\bar{\mathbf{w}}^{j-1} - \bar{X}\hat{\mathbf{w}}^{j-1}\| \\ &\leq 2\|\boldsymbol{\alpha}^j - \boldsymbol{\alpha}^{j-1}\| + \|\bar{X}\bar{\mathbf{w}}^j - \bar{X}\bar{\mathbf{w}}^{j-1}\| + \|\bar{X}\hat{\mathbf{w}}^j - \bar{X}\bar{\mathbf{w}}^j\| + \|\bar{X}\bar{\mathbf{w}}^{j-1} - \bar{X}\hat{\mathbf{w}}^{j-1}\| \\ &\leq (2 + M)\|\boldsymbol{\alpha}^j - \boldsymbol{\alpha}^{j-1}\| + \sum_{t=j-\tau}^{j-1} \|\Delta\alpha_t\|M + \sum_{t=j-\tau-1}^{j-2} \|\Delta\alpha_t\|M \\ &= (2 + 2M)\|\boldsymbol{\alpha}^j - \boldsymbol{\alpha}^{j-1}\| + 2M \sum_{t=j-\tau-1}^{j-2} \|\Delta\alpha_t\| \end{aligned} \quad (22)$$

Now we prove (8) by induction.

**Induction Hypothesis.** Due to Proposition 1, we prove the following equivalent statement. For all  $j$ ,

$$E(\|\boldsymbol{\alpha}^{j-1} - \tilde{\boldsymbol{\alpha}}^j\|^2) \leq \rho E(\|\boldsymbol{\alpha}^j - \tilde{\boldsymbol{\alpha}}^{j+1}\|^2), \quad (23)$$

**Induction Basis.** When  $j = 1$ ,

$$\|\boldsymbol{\alpha}^1 - \tilde{\boldsymbol{\alpha}}^2 + \boldsymbol{\alpha}^0 - \tilde{\boldsymbol{\alpha}}^1\| \leq (2 + 2M)\|\boldsymbol{\alpha}^1 - \boldsymbol{\alpha}^0\|.$$

By taking the expectation on (21), we have

$$\begin{aligned} E[\|\alpha^0 - \tilde{\alpha}^1\|^2] - E[\|\alpha^1 - \tilde{\alpha}^2\|^2] &\leq 2E[\|\alpha^0 - \tilde{\alpha}^1\| \|\alpha^1 - \tilde{\alpha}^2 - \alpha^0 + \tilde{\alpha}^1\|] \\ &\leq (4 + 4M)E(\|\alpha^0 - \tilde{\alpha}^1\| \|\alpha^0 - \alpha^1\|). \end{aligned}$$

From (17) we have  $E[\|\alpha^0 - \alpha^1\|^2] = \frac{1}{n}\|\alpha^0 - \tilde{\alpha}^1\|^2$ . Also, by AM-GM inequality, for any  $\mu_1, \mu_2 > 0$  and any  $c > 0$ , we have

$$\mu_1 \mu_2 \leq \frac{1}{2}(c\mu_1^2 + c^{-1}\mu_2^2). \quad (24)$$

Therefore, we have

$$\begin{aligned} &E[\|\alpha^0 - \tilde{\alpha}^1\| \|\alpha^0 - \alpha^1\|] \\ &\leq \frac{1}{2}E\left[n^{1/2}\|\alpha^0 - \alpha^1\|^2 + n^{-1/2}\|\tilde{\alpha}^1 - \alpha^0\|^2\right] \\ &= \frac{1}{2}E\left[n^{-1/2}\|\alpha^0 - \tilde{\alpha}^1\|^2 + n^{-1/2}\|\tilde{\alpha}^1 - \alpha^0\|^2\right] \quad \text{by (17)} \\ &= n^{-1/2}E[\|\alpha^0 - \tilde{\alpha}^1\|^2]. \end{aligned}$$

Therefore,

$$E[\|\alpha^0 - \tilde{\alpha}^1\|^2] - E[\|\alpha^1 - \tilde{\alpha}^2\|^2] \leq \frac{4 + 4M}{\sqrt{n}}E[\|\alpha^0 - \tilde{\alpha}^1\|^2],$$

which implies

$$E[\|\alpha^0 - \alpha^1\|^2] \leq \frac{1}{1 - \frac{4+4M}{\sqrt{n}}}E[\|\alpha^1 - \tilde{\alpha}^2\|^2] \leq \rho E[\|\alpha^1 - \tilde{\alpha}^2\|^2], \quad (25)$$

where the last inequality is based on Proposition 4 and the fact  $\theta M \geq 1$ .

**Induction Step.** By the induction hypothesis, we assume

$$E[\|\alpha^{t-1} - \tilde{\alpha}^t\|^2] \leq \rho E[\|\alpha^t - \tilde{\alpha}^{t+1}\|^2] \quad \forall t \leq j-1. \quad (26)$$

The goal is to show

$$E[\|\alpha^{j-1} - \tilde{\alpha}^j\|^2] \leq \rho E[\|\alpha^j - \tilde{\alpha}^{j+1}\|^2].$$

First, we show that for all  $t < j$ ,

$$E\left[\|\alpha^t - \alpha^{t+1}\| \|\alpha^{j-1} - \tilde{\alpha}^j\|\right] \leq \frac{\rho^{(j-1-t)/2}}{\sqrt{n}}E\left[\|\alpha^{j-1} - \tilde{\alpha}^j\|^2\right] \quad (27)$$

*Proof.* By (24) with  $c = n^{1/2}\beta$ , where  $\beta = \rho^{(t+1-j)/2}$ ,

$$\begin{aligned}
& E\left[\|\alpha^t - \alpha^{t+1}\| \|\alpha^{j-1} - \tilde{\alpha}^j\|\right] \\
& \leq \frac{1}{2}E\left[n^{1/2}\beta\|\alpha^t - \alpha^{t+1}\|^2 + n^{-1/2}\beta^{-1}\|\alpha^{j-1} - \tilde{\alpha}^j\|^2\right] \\
& = \frac{1}{2}E\left[n^{1/2}\beta E[\|\alpha^t - \alpha^{t+1}\|^2] + n^{-1/2}\beta^{-1}\|\alpha^{j-1} - \tilde{\alpha}^j\|^2\right] \\
& = \frac{1}{2}E\left[n^{-1/2}\beta\|\alpha^t - \tilde{\alpha}^{t+1}\|^2 + n^{-1/2}\beta^{-1}\|\alpha^{j-1} - \tilde{\alpha}^j\|^2\right] \quad \text{by Proposition 1} \\
& \leq \frac{1}{2}E\left[n^{-1/2}\beta\rho^{j-1-t}\|\alpha^{j-1} - \tilde{\alpha}^j\|^2 + n^{-1/2}\beta^{-1}\|\alpha^{j-1} - \tilde{\alpha}^j\|^2\right] \quad \text{by Eq. (26)} \\
& \leq \frac{1}{2}E\left[n^{-1/2}\beta^{-1}\|\alpha^{j-1} - \tilde{\alpha}^j\|^2 + n^{-1/2}\beta^{-1}\|\tilde{\alpha}^{j-1} - \alpha^j\|^2\right] \quad \text{by the definition of } \beta \\
& \leq \frac{\rho^{(j-1-t)/2}}{\sqrt{n}}E\left[\|\alpha^{j-1} - \tilde{\alpha}^j\|^2\right]
\end{aligned}$$

□

Let  $\theta = \sum_{t=1}^{\tau} \rho^{t/2}$ . We have

$$\begin{aligned}
& E[\|\alpha^{j-1} - \tilde{\alpha}^j\|^2] - E[\|\alpha^j - \tilde{\alpha}^{j+1}\|^2] \\
& \leq E\left[2\|\alpha^{j-1} - \tilde{\alpha}^j\|((2+2M)\|\alpha^j - \alpha^{j-1}\| + 2M \sum_{t=j-\tau-1}^{j-1} \|\alpha^t - \alpha^{t-1}\|)\right] \quad \text{by (21), (22)} \\
& = (4+4M)E(\|\alpha^{j-1} - \tilde{\alpha}^j\| \|\alpha^j - \alpha^{j-1}\|) + 4M \sum_{t=j-\tau-1}^{j-1} E\left[\|\alpha^{j-1} - \tilde{\alpha}^j\| \|\alpha^t - \alpha^{t-1}\|\right] \\
& \leq (4+4M)n^{-1/2}E[\|\tilde{\alpha}^j - \alpha^{j-1}\|^2] + 4Mn^{-1/2}E[\|\alpha^{j-1} - \tilde{\alpha}^j\|^2] \sum_{t=j-1-\tau}^{j-2} \rho^{(j-1-t)/2} \quad \text{by (27)} \\
& \leq (4+4M)n^{-1/2}E[\|\tilde{\alpha}^j - \alpha^{j-1}\|^2] + 4Mn^{-1/2}\theta E[\|\alpha^{j-1} - \tilde{\alpha}^j\|^2] \\
& \leq \frac{4+4M+4M\theta}{\sqrt{n}}E[\|\alpha^{j-1} - \tilde{\alpha}^j\|^2],
\end{aligned}$$

which implies that

$$E[\|\alpha^{j-1} - \tilde{\alpha}^j\|^2] \leq \frac{1}{1 - \frac{4+4M+4M\theta}{\sqrt{n}}}E[\|\alpha^j - \tilde{\alpha}^{j+1}\|^2] \leq \rho E[\|\alpha^j - \tilde{\alpha}^{j+1}\|^2],$$

where the last inequality is based on Proposition 4.

### A.3 Proof of Theorem 2

First, we define  $T(\mathbf{w}, \alpha)$  to be a  $n$ -dimensional vector such that

$$(T(\mathbf{w}, \alpha))_t = T_t(\mathbf{w}, \alpha_t) \text{ for all } t,$$

We can then bound the distance  $E[\|T(\mathbf{w}^j, \boldsymbol{\alpha}^j) - T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j)\|^2]$  by (we omit the expectation in the following derivation):

$$\begin{aligned}
\|T(\mathbf{w}^j, \boldsymbol{\alpha}^j) - T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j)\|^2 &= \sum_{t=1}^n (T_t(\mathbf{w}^j, \boldsymbol{\alpha}_t^j) - T_t(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}_t^j))^2 \\
&\leq \sum_t \left( \frac{(\mathbf{w}^j - \hat{\mathbf{w}}^j)^T \mathbf{x}_t}{\|\mathbf{x}_t\|^2} \right)^2 \quad (\text{By Proposition 2}) \\
&= \|\bar{X}(\mathbf{w}^j - \hat{\mathbf{w}}^j)\|^2 \\
&\leq M^2 \left( \sum_{t=j-\tau}^{j-1} \|\boldsymbol{\alpha}^{t+1} - \boldsymbol{\alpha}^t\| \right)^2 \quad (\text{By Proposition 3}) \\
&\leq \tau M^2 \left( \sum_{t=j-\tau}^{j-1} \|\boldsymbol{\alpha}^{t+1} - \boldsymbol{\alpha}^t\|^2 \right) \\
&\leq \tau M^2 \left( \sum_{t=1}^{\tau} \rho^t \|\boldsymbol{\alpha}^j - \boldsymbol{\alpha}^{j+1}\|^2 \right) \quad (\text{By Lemma 1}) \\
&\leq \frac{\tau M^2}{n} \left( \sum_{t=1}^{\tau} \rho^t \right) \|T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2 \\
&\leq \frac{\tau^2 M^2}{n} \rho^\tau \|T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2
\end{aligned}$$

Since  $\rho^{(\tau+1)/2} \leq e$ , we have  $\rho^{\tau+1} \leq e^2$ , so  $\rho^\tau \leq e^2$  since  $\rho \geq 1$ . Therefore,

$$\|T(\mathbf{w}^j, \boldsymbol{\alpha}^j) - T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j)\|^2 \leq \frac{\tau^2 M^2 e^2}{n} \|T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2. \quad (28)$$

As a result,

$$\begin{aligned}
\|T(\mathbf{w}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2 &= \|T(\mathbf{w}^j, \boldsymbol{\alpha}^j) - T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j) + T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2 \\
&\leq 2(\|T(\mathbf{w}^j, \boldsymbol{\alpha}^j) - T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j)\|^2 + \|T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2) \\
&\leq 2\left(1 + \frac{e^2 \tau^2 M^2}{n}\right) \|T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2. \quad (29)
\end{aligned}$$

Next, we bound the decrease of objective function value by

$$\begin{aligned}
D(\boldsymbol{\alpha}^j) - D(\boldsymbol{\alpha}^{j+1}) &= D(\boldsymbol{\alpha}^j) - D(\bar{\boldsymbol{\alpha}}^{j+1}) + D(\bar{\boldsymbol{\alpha}}^{j+1}) - D(\boldsymbol{\alpha}^{j+1}) \\
&\geq \frac{\|\mathbf{x}_{i(j)}\|^2}{2} \|\boldsymbol{\alpha}_{i(j)}^j - T_{i(j)}(\mathbf{w}^j, \boldsymbol{\alpha}^j)\|^2 - \frac{L_{max}}{2} \|T_{i(j)}(\mathbf{w}^j, \boldsymbol{\alpha}^j) - T_{i(j)}(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j)\|^2
\end{aligned}$$



So

$$\begin{aligned}
E[D(\boldsymbol{\alpha}^j)] - E[D(\boldsymbol{\alpha}^{j+1})] &\geq \frac{R_{min}^2}{2n} E[\|T(\mathbf{w}^j, \boldsymbol{\alpha}^j)\|^2] - \frac{L_{max}}{2n} E[\|T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2] \\
&\geq \frac{R_{min}^2}{2n} E[\|T(\mathbf{w}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2] - \frac{L_{max}}{2n} \frac{\tau^2 M^2 e^2}{n} E[\|T(\hat{\mathbf{w}}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2] \\
&\geq \frac{R_{min}^2}{2n} E[\|T(\mathbf{w}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2] - \frac{2L_{max}}{2n} \frac{\tau^2 M^2 e^2}{n} (1 + \frac{e\tau M}{\sqrt{n}}) E[\|T(\mathbf{w}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2] \\
&\geq \frac{R_{min}^2}{2n} \left(1 - \frac{2L_{max}}{R_{min}^2} (1 + \frac{e\tau M}{\sqrt{n}}) (\frac{\tau^2 M^2 e^2}{n})\right) E[\|T(\mathbf{w}^j, \boldsymbol{\alpha}^j) - \boldsymbol{\alpha}^j\|^2]
\end{aligned}$$

Let  $b = (1 - \frac{2L_{max}}{R_{min}^2} (1 + \frac{e\tau M}{\sqrt{n}}) (\frac{\tau^2 M^2 e^2}{n}))$  and combine the above inequality with eq (9) we have

$$\begin{aligned}
E[D(\boldsymbol{\alpha}^j)] - E[D(\boldsymbol{\alpha}^{j+1})] &\geq b\kappa E[\|\boldsymbol{\alpha}^j - P_S(\boldsymbol{\alpha}^j)\|^2] \\
&\geq \frac{b\kappa}{L_{max}} E[D(\boldsymbol{\alpha}^j) - D^*].
\end{aligned}$$

Therefore, we have

$$\begin{aligned}
E[D(\boldsymbol{\alpha}^{j+1})] - D^* &= E[D(\boldsymbol{\alpha}^j)] - (E[D(\boldsymbol{\alpha}^j)] - E[D(\boldsymbol{\alpha}^{j+1})]) - D^* \\
&\leq (1 - \frac{b\kappa}{L_{max}}) (E[D(\boldsymbol{\alpha}^j)] - D^*).
\end{aligned}$$